

# A LEVEL-SET HIT-AND-RUN SAMPLER FOR QUASI-CONCAVE DISTRIBUTIONS

Dean Foster and Shane T. Jensen

Department of Statistics  
The Wharton School  
University of Pennsylvania  
stjensen@wharton.upenn.edu

## Abstract

We develop a new sampling strategy that uses the hit-and-run algorithm within level sets of the target density. Our method can be applied to any quasi-concave density, which covers a broad class of models. Our sampler performs well in high-dimensional settings, which we illustrate with a comparison to Gibbs sampling on a spike-and-slab mixture model. We also extend our method to exponentially-tilted quasi-concave densities, which arise often in Bayesian models consisting of a log-concave likelihood and quasi-concave prior density. Within this class of models, our method is effective at sampling from posterior distributions with high dependence between parameters, which we illustrate with a simple multivariate normal example. We also implement our level-set sampler on a Cauchy-normal model where we demonstrate the ability of our level set sampler to handle multi-modal posterior distributions.

February 21, 2012

# 1 Introduction

Complex statistical models are often estimated by sampling random variables from complicated distributions. This strategy is especially prevalent within the Bayesian framework, where Markov Chain Monte Carlo (MCMC) simulation is used to estimate the posterior distribution of a set of unknown parameters. The most common MCMC technique is the Gibbs sampler (Geman and Geman, 1984), where small subsets (often one parameter at a time) are sampled from their conditional posterior distribution given the current values of all other parameters. This component-wise strategy reduces the sampling from complicated multi-dimension distributions into sampling from a series of small (or one-) dimension distributions. For sampling in low dimensions, various techniques such as the Metropolis-Hastings algorithm (Hastings, 1970) can be employed.

However, in high dimensional parameter spaces component-wise strategies can encounter problems such as high autocorrelation and the inability to move between local modes. In high dimensions, one would ideally employ a strategy that sampled the high dimensional parameter vector directly, instead of one component at a time. In this paper, we develop a sampling algorithm that provides direct samples from a (potentially high-dimensional) *quasi-concave* density function.

Consider a high-dimensional density function  $f(\mathbf{x})$  from which we would like to obtain samples  $\mathbf{x}$ . A density function  $f$  is quasi-concave if:

$$C_a = \{\mathbf{x} \mid f(\mathbf{x}) > a\}$$

is convex for all values of  $a$ . Our procedure is based on the fact that any hori-

zontal slice through a quasi-concave density  $f$  will give a convex level set above that slice. By slicing the quasi-concave density  $f$  at a sequence of different heights, we divide the density  $f$  into a sequence of convex level sets.

We then use the hit-and-run algorithm to sample high-dimensional points  $x$  within each of these convex level sets, while simultaneously estimating the volume of each convex level set. As reviewed in Vempala (2005), recent work suggests that the hit-and-run algorithm is an efficient way to sample from high-dimension convex sets as long as the start is “warm”. The hit-and-run algorithm is not as commonplace as other sampling methods (e.g. Metropolis-Hastings), though Chen and Schmeiser (1996) discuss using hit-and-run Monte Carlo to evaluate multi-dimensional integrals.

In Section 2.1, we review several recent results for hit-and-run sampling in convex sets. In Section 2.2, we outline our procedure for ensuring a warm start within each convex slice, thereby giving us an efficient way of sampling from the entire quasi-concave density. In Section 3, we present an empirical comparison that suggests our “level-set hit-and-run sampling” methodology is much more efficient than component-wise sampling schemes for high dimensional quasi-concave densities.

We also extend our method to sample efficiently from *exponentially-tilted quasi-concave* densities in Section 2.4. Exponentially-tilted quasi-concave densities are very common in Bayesian models: the combination of a quasi-concave prior density and a log-concave likelihood leads to an exponentially-tilted quasi-concave posterior density.

In Section 4, we implement our exponentially-tilted level-set sampler on a simple multivariate normal density. In high dimensions, our method is more ef-

fective at obtaining posterior samples than a component-wise Gibbs sampling strategy.

Finally, in Section 5, we illustrate the efficiency of our method on one such model, consisting of a Normal data likelihood and a Cauchy prior density. It should be noted that this posterior density can be multi-modal, depending on the relative locations of the prior mode and observed data. Our results in Section 5 suggest that our extended method can accommodate posterior distributions which are both high-dimensional and multi-modal.

## 2 Methodology and Theory

### 2.1 Quasi-Concave Densities and Sampling Convex Level Sets

Let  $f()$  be a density in a high dimensional space, then  $f$  is quasi-concave if:

$$C_a = \{\mathbf{x} \mid f(\mathbf{x}) > a\}$$

is convex for all values of  $a$ . In other words, the level set  $C_a$  of a quasi-concave density  $f$  is convex for any value  $a$ . Let  $\mathcal{Q}$  denote the set of all quasi-concave densities and  $\mathcal{D}$  denote the set of all concave densities. All concave densities are quasi-concave,  $\mathcal{D} \subset \mathcal{Q}$ , but the converse is not true. Quasi-concave densities are a very broad class that contains many common data models, including the normal, gamma, student's  $t$ , and uniform densities.

Our proposed methodology is based on computing the volume of a convex level set  $C$  by obtaining random samples  $\mathbf{x}$  from  $C$ . Let us assume we are already at a point  $\mathbf{x}^0$  in  $C_a$ . A simple algorithm for obtaining a random sample

$\mathbf{x}$  would be a *ball walk*: pick a uniform point  $\mathbf{x}$  in a ball of size  $\delta$  around  $\mathbf{x}^0$ , and move to  $\mathbf{x}$  only if  $\mathbf{x}$  is still in  $C$ .

An alternative algorithm is *hit-and-run*, where a random direction  $\mathbf{d}$  is picked from current point  $\mathbf{x}^0$ . This direction will intersect the boundary of the convex set  $C$  at some point  $\mathbf{x}^1$ . A new point  $\mathbf{x}$  is sampled uniformly along the line segment defined by direction  $\mathbf{d}$  and end points  $\mathbf{x}^0$  and  $\mathbf{x}^1$ , and is thus guaranteed to remain in  $C$ .

Lovasz (1999) showed that the hit-and-run algorithm mixes rapidly from a *warm start* in a convex body. The warm start criterion is designed to ensure that our starting point is not stuck in some isolated corner of the convex body.

Vempala (2005) suggests that a random point from convex set  $C' \in C$  provides a warm start for convex set  $C$  as long as  $\text{volume}(C')/\text{volume}(C) \geq 0.5$ . Vempala (2005) also presents several results that suggest the hit-and-run algorithm mixes more rapidly than the ball walk algorithm.

## 2.2 LSHR1: Level-Set Hit-and-run Sampler for Quasi-concave Densities

We harness the results of the previous section in our sampling algorithm for an arbitrary quasi-concave probability density. Briefly, our algorithm proceeds in stages where we sample from increasingly larger level sets of the quasi-concave density, while ensuring that each level set provides a warm start for the next level set. After the entire density has been covered in this fashion, our samples from each level set are weighted appropriately to provide a full sample from the quasi-concave probability density. We will see in later sections that our

procedure is far more efficient than standard MCMC methods when the desired quasi-concave density has high dimension. We outline our algorithm in more detail below, and also include pseudo-code in Appendix A.

Our algorithm starts by taking a small level set centered around the maximum value  $\mathbf{x}_{\max}$  of the quasi-concave density  $f(\cdot)$ . This initial slice is defined as the convex shape  $C_1$  of the probability density  $f(\cdot)$  above an initial density threshold  $t_1$ . In other words, a particular point  $\mathbf{x}$  is a valid member of slice  $C_1$  iff  $f(\mathbf{x}) \geq t_1$ .

Starting from  $\mathbf{x}_{\max}$ , we run a hit-and-run sampler within this initial level set for  $m$  iterations. In our supplementary materials, we explore appropriate choices for the number of hit-and-run iterations  $m$ .

Each iteration of the hit-and-run sampler consists of first picking a random direction<sup>1</sup>  $\mathbf{d}$  from current point  $\mathbf{x}$ . This direction  $\mathbf{d}$  and current point  $\mathbf{x}$  define a line segment with endpoints at the edge of the convex level set  $C_1$ . A new point  $\mathbf{x}'$  is sampled uniformly along this line segment, which ensures that  $\mathbf{x}'$  remains in  $C_1$ . We use  $\{\mathbf{x}\}_1$  to denote the set of all points sampled from  $C_1$  using this hit-and-run method.

After sampling  $m$  times from convex level set  $C_1$ , we need to pick a new threshold  $t_2 < t_1$  that specifies a new level set of the quasi-concave probability density  $f(\cdot)$ . However, in order for our sampler to stay “warm” as defined in Section 2.1, we need  $t_2$  to define a convex shape  $C_2$  with volume  $V_2$  that is less than twice the volume of our initial volume  $V_1$ .

We define the volume ratio of two level sets as  $R_{1:2} = V_1/V_2$ , so we need  $R_{1:2} \geq$

---

<sup>1</sup>Note that we can increase the efficiency of the hit-and-run sampler for highly non-spherical convex shapes by scaling this random direction by the estimated covariance matrix of our samples. See Appendix A for details.

0.5 for the two level sets  $C_1$  and  $C_2$ . In addition, we want a threshold  $t_2$  that is not too close to  $t_1$  so that we are not using an excessive number of level sets.

We meet these criteria by first proposing a new threshold  $t_{\text{prop}}$  and then, starting from current point  $\mathbf{x}$ , running  $m$  iterations of the hit-and-run sampler within the new convex level set  $C_{\text{prop}}$  defined by  $f(\mathbf{x}) > t_{\text{prop}}$ . We focus on the ratio of volumes  $R_{1:\text{prop}} = V_1/V_{\text{prop}}$  which we estimate with  $\hat{R}_{\text{prop}}$  equal to the proportion of these sampled points from convex level set  $C_{\text{prop}}$  that were also contained within the previous level set  $C_1$ .

We only accept the proposed threshold  $t_{\text{prop}}$  if  $0.55 \leq \hat{R}_{\text{prop}} \leq 0.8$ . The lower bound fulfills the criterion for the new level to be “warm” and the *ad hoc* upper bound fulfills our desire for efficiency: we do not want the new level set to not be too close in volume to the previous level set.

If the proposed threshold is not accepted, we propose a new threshold  $t_{\text{prop}}$  in an adaptive fashion (details in Appendix A) and repeat the same steps for testing this new threshold. If the proposed threshold is accepted, we re-define  $t_{\text{prop}} \equiv t_2$  and then the convex level set  $C_{\text{prop}} \equiv C_2$  becomes our current level set. The collection of sampled points  $\{\mathbf{x}\}_2$  from  $C_2$  are retained, as is our estimated ratio of slice volumes  $\hat{R}_{1:2}$ .

In this same way, we continue to add level sets  $C_k$  defined by threshold  $t_k$  based on comparison to the previous level sets  $C_{k-1}$ . The level set algorithm terminates when the accepted threshold  $t_k$  is lower than some pre-specified lower bound  $K$  on the probability density  $f(\cdot)$ .

The result of this entire procedure is  $n$  level sets, represented by a vector of thresholds  $(t_1, \dots, t_n)$ , a vector of estimated ratios of volumes  $(\hat{R}_{1:2}, \dots, \hat{R}_{n-1:n})$ , and the level set collections: a set of  $m$  sampled points from each level set

$(\{\mathbf{x}\}_1, \dots, \{\mathbf{x}\}_n).$

In order to obtain  $L$  essentially independent<sup>2</sup> samples from the pseudo-convex density, we sub-sample points from the level set collections  $(\{\mathbf{x}\}_1, \dots, \{\mathbf{x}\}_n)$ , with each level set  $i$  represented proportional to its probability  $p_i$ .

The probability of each level set is the product of its density height and volume,

$$p_i = (t_{i-1} - t_i) \times V_i$$

We have our volume ratios  $\hat{R}_{i:i+1}$  as estimates of  $V_i/V_{i+1}$ , which allows us to estimate  $p_i$  by first calculating

$$\hat{q}_i = (t_{i-1} - t_i) \times \prod_{j=i}^n \hat{R}_{j:j+1} \quad i = 1, \dots, n \quad (1)$$

where  $t_0$  is the maximum of  $f(\cdot)$  and  $\hat{R}_{n:n+1} = 1$ , and then  $\hat{p}_i = \hat{q}_i / \sum \hat{q}_j$ .

We demonstrate our LSHR1 algorithm on a spike-and-slab density in Section 3.

## 2.3 Exponentially-Tilted Quasi-concave Densities

A density  $g()$  is a log-concave density function if  $\log g()$  is a concave density function. Let  $\mathcal{L}$  denote the set of all log-concave density functions. All log-concave density functions are also concave density functions ( $\mathcal{L} \subset \mathcal{D}$ ), and thus all log-concave density functions are also quasi-concave density functions ( $\mathcal{L} \subset \mathcal{Q}$ ).

---

<sup>2</sup>Although there is technically some dependence between successive sampled points from the hit-and-run sampler, our scheme of sub-sampling randomly from level set collections  $(\{\mathbf{x}\}_1, \dots, \{\mathbf{x}\}_n)$  essentially removes this dependence



Bagnoli and Bergstrom (2005) gives an excellent review of log-concave densities. The normal density is log-concave whereas the student's  $t$  and Cauchy density are not. The gamma and beta densities are log-concave only under certain parameter settings, e.g. both the uniform and exponential densities are log-concave. The beta( $a, b$ ) density with other parameter settings (e.g.  $a < 1, b < 1$ ) can be neither log-concave nor quasi-concave.

Now, let  $\mathcal{T}$  denote the set of *exponentially-tilted* quasi-concave density functions. A density function  $h$  is an exponentially-tilted quasi-concave density function if there exists a quasi-concave density  $f \in \mathcal{Q}$  such that  $f(\mathbf{x})/h(\mathbf{x}) = \exp(\boldsymbol{\beta}'\mathbf{x})$ . Exponentially-tilted quasi-concave densities are a generalization of quasi-concave densities, so  $\mathcal{Q} \subset \mathcal{T}$ .

These three classes of functions (log-concave, quasi-concave, and exponentially-tilted quasi-concave) are linked by the following important relationship: if  $X \sim f$  where  $f \in \mathcal{Q}$  and  $Y|X \sim g$  where  $g \in \mathcal{L}$ , then  $X|Y \sim h$  where  $h \in \mathcal{T}$ .

The consequences of this relationship is apparent for the Bayesian approach. The combination of a quasi-concave prior density for parameters  $\boldsymbol{\theta}$  and a log-concave likelihood for data  $\mathbf{y}|\boldsymbol{\theta}$  will produce an exponentially-tilted quasi-concave posterior density  $\boldsymbol{\theta}|\mathbf{y}$ .

We now extend our level-set hit-and-run sampling methodology to exponentially-tilted quasi-concave densities, which makes our procedure applicable to a large class of Bayesian models consisting of quasi-concave priors and log-concave likelihoods.

As mentioned above, examples of log-concave likelihoods for data  $\mathbf{y}|\boldsymbol{\theta}$  include the normal density, exponential density and uniform density. Quasi-concave priors for parameters  $\boldsymbol{\theta}$  are an even broader class of densities, including the

normal density, the  $t$  density, the Cauchy density and the gamma density.

In Section 4, we examine a simple multivariate normal density, which is log-concave. Then, in Section 5, we examine an exponentially-tilted quasi-concave posterior density resulting from the combination of a multivariate normal density with a Cauchy prior distribution.

## 2.4 LSHR2: Level-Set Hit-and-Run Sampler for Exponentially-Tilted Quasi-concave Densities

In Section 2.2, we presented our level set hit-and-run algorithm for sampling  $\mathbf{x}$  from quasiconcave density  $f(\mathbf{x})$ . We now extend our algorithm to sample from an exponentially-tilted quasi-concave density  $h(\mathbf{x})$ .

As mentioned in Section 2.3, exponentially-tilted quasi-concave densities commonly arise as posterior distributions in Bayesian models. In keeping with the usual notation for Bayesian models, we replace our previous variables  $\mathbf{x}$  with parameters  $\boldsymbol{\theta}$ . These parameters  $\boldsymbol{\theta}$  have an exponentially-tilted quasi-concave posterior density  $h(\boldsymbol{\theta}|\mathbf{y})$  arising from a log-concave likelihood  $g(\mathbf{y}|\boldsymbol{\theta})$  and quasi-concave prior density  $f(\boldsymbol{\theta})$ .

Our LSHR2 algorithm starts just as before, by taking a small level set centered around the maximum value  $\boldsymbol{\theta}_{\max}$  of the quasi-concave prior density  $f(\cdot)$ . This initial level set is defined, in part, as the convex shape of the probability density  $f(\cdot)$  above an initial density threshold  $t_1$ .

However, we now augment our sampling space with an extra variable  $p$  where

$p < g(\mathbf{y}|\log \boldsymbol{\theta})$ . Letting  $\boldsymbol{\theta}^* = (\boldsymbol{\theta}, p)$ , our convex shape is now

$$D_1 = \{\boldsymbol{\theta}^* : f(\boldsymbol{\theta}) \geq t_1, p < \log g(\mathbf{y}|\boldsymbol{\theta})\}$$

Within this new convex shape, we run an exponentially-weighted version of the hit-and-run algorithm. Specifically, a random  $(d+1)$ -dimensional direction  $\mathbf{d}$  is sampled<sup>3</sup> which, along with the current  $\boldsymbol{\theta}^*$ , defines a line segment with endpoints at the boundaries of  $D_1$ . Instead of sampling uniformly along this line segment, we sample a new point  $(\boldsymbol{\theta}^*)'$  from points on this line segment weighted by  $\exp(p)$ .

The remainder of the LSHR2 algorithm proceeds in the same fashion as our LSHR1 algorithm: we construct a schedule of decreasing thresholds  $t_1, t_2, \dots$  and corresponding level sets  $D_1, D_2, \dots$  such that each level set  $D_k$  is a warm start for the next level set  $D_{k+1}$ .

Within each of these steps, the exponentially-weighted hit-and-run algorithm is used to sample values  $\boldsymbol{\theta}^*$  within the current level set. As before,  $D_k$  is a warm start for the next level set  $D_{k+1}$  if the ratio of volumes<sup>4</sup>  $\hat{R}_{k:k+1} = V_k/V_{k+1}$  that lies within the range  $0.55 \leq \hat{R}_{k:k+1} \leq 0.8$ . The algorithm terminates when our decreasing thresholds  $t_k$  achieve a pre-specified lower bound  $K$ .

Our procedure results in  $n$  level sets  $(D_1, \dots, D_n)$ , represented by a vector of thresholds  $(t_1, \dots, t_n)$ , a vector of estimated ratios of volumes  $(\hat{R}_{1:2}, \dots, \hat{R}_{n-1:n})$ ,

---

<sup>3</sup>Again we can increase the efficiency of the hit-and-run sampler for highly non-spherical convex shapes by scaling this random direction by the estimated covariance matrix of our samples. See Appendix B for details.

<sup>4</sup>We continue to use the term “volume” even though these convex shapes are a combination of  $d$ -dimensional  $\boldsymbol{\theta}$  and the extra log-density  $p$  dimension.

and the level set collections: a set of  $m$  sampled points from each level set:

$$(\{\boldsymbol{\theta}^*\}_1, \dots, \{\boldsymbol{\theta}^*\}_n)$$

Finally, we obtain  $L$  samples  $(\boldsymbol{\theta}_1^*, \dots, \boldsymbol{\theta}_L^*)$  by sub-sampling points from the level set collections  $(\{\boldsymbol{\theta}^*\}_1, \dots, \{\boldsymbol{\theta}^*\}_n)$ , with each level set  $i$  represented proportional to its probability  $p_i$ . The probability of each level set is still calculated as in (1).

By simply ignoring the  $(d+1)$ -th sampled dimension  $p$ , we have samples  $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_L)$  from the exponentially-tilted pseudo-convex posterior density  $h(\boldsymbol{\theta}|\mathbf{y})$ . Details of our LSHR2 algorithm are given in Appendix B. We demonstrate our LSHR2 algorithm on a multivariate normal density in Section 4 and a Cauchy-normal model in Section 5.

## 2.5 Component-wise Alternatives: Gibbs and Metropolis-Hastings

Although there have been many recent advances in MCMC methods, we focus on our comparisons on the Gibbs sampler as it remains the most popular method for obtaining samples from high-dimensional distributions. The Gibbs sampler (Geman and Geman, 1984) is a Markov Chain Monte Carlo approach to sampling from a multi-dimensional density  $f(\mathbf{x})$ .

The basic Gibbs strategy is to sample a single component  $x_i$  from its univariate conditional distribution  $f(x_i|\mathbf{x}_{-i})$  given all other components  $\mathbf{x}_{-i}$ . Successively iterating this sampling strategy through each component  $i = 1, \dots, n$  will give a full sample  $\mathbf{x}$ . These sampled  $\mathbf{x}$ 's will converge in distribution to the target distribution  $f(\mathbf{x})$ .

The benefit of a Gibbs approach is that sampling from a high-dimensional den-

sity  $f(\mathbf{x})$  is reduced to sampling from a series of one-dimensional densities  $f(x_i|\mathbf{x}_{-i})$ . The Gibbs sampler is most attractive when these one-dimensional conditional distributions each have a standard form but even if this is not the case, the Metropolis-Hastings algorithm (Hastings, 1970) can be employed for sampling from non-standard distributions.


The Gibbs sampler is not always strictly component-wise: for some models, blocks of variables can be sampled in a single step from a standard distribution. That said, a Gibbs strategy for the sampling of a high-dimensional  $\mathbf{x}$  will typically involve sampling a series of small subsets of  $\mathbf{x}$ .

The component-wise approach of the Gibbs sampler has both strengths and weaknesses. The relative simplicity of sampling from conditional distributions is a clear strength. However, a potential weakness is that high correlation can lead to extremely slow convergence to the target distribution. Our own approach is motivated by the belief that it is advantageous to sample the entire set of variables simultaneously when possible.

### 3 Empirical Study 1: Spike-and-slab Density

We consider the problem of sampling from a multivariate density which consists of a 50-50 mixture of two normal distributions, both centered at zero but different variances. Specifically, the first component has variance  $\Sigma_0$  with off-diagonal elements equal to zero and diagonal elements equal to 0.05, whereas the second component has variance  $\Sigma_0$  with off-diagonal elements equal to zero and diagonal elements equal to 3. Figure 1 gives this spike-and-slab density in a single dimension. The spike-and-slab density is commonly used as a mixture model for important versus spurious predictors in variable selection models

**Figure 1:** Spike-and-slab density: gray lines indicate density of the two components, black line is the mixture density



(George and McCulloch, 1997).

This density is quasi-concave and thus amenable to our proposed level-set hit-and-run sampling methodology. Sampling from this density in a single dimension can be easily accomplished by grid sampling or Gibbs sampling using data augmentation. However, we will see that these simpler strategies will not perform adequately in higher dimensions whereas our LSHR1 sampling strategy continues to show good performance in high dimensions.

### 3.1 Gibbs Sampling for Spike and Slab

The usual Gibbs sampling approach to obtaining a sample  $\mathbf{x}$  from a mixture density is to augment the variable space with an indicator variable  $I$  where

$I = 1$  indicates the current mixture component from which  $x$  is drawn. The algorithm iterates between

1. Sample  $x$  from mixture component dictated by  $I$ , i.e.

$$\mathbf{x} \sim \begin{cases} \mathcal{N}(\mathbf{0}, \Sigma_1) & \text{if } I = 1 \\ \mathcal{N}(\mathbf{0}, \Sigma_0) & \text{if } I = 0 \end{cases}$$

2. Sample  $I$  with probability based on  $x$ :  $P(I = 1) = \frac{\phi(\mathbf{x}, \mathbf{0}, \Sigma_1)}{\phi(\mathbf{x}, \mathbf{0}, \Sigma_1) + \phi(\mathbf{x}, \mathbf{0}, \Sigma_0)}$


where  $\phi(\mathbf{x}, \boldsymbol{\mu}, \Sigma)$  is the density of a multivariate normal with mean  $\boldsymbol{\mu}$  and variance  $\Sigma$  evaluated at  $\mathbf{x}$ .

This algorithm mixes well when  $\mathbf{x}$  has a small number of dimensions ( $d < 10$ ), but in higher dimensions, it is difficult for the algorithm to move between the two components. We see this behavior in Figure 2, where we evaluate results from running the Gibbs sampler for 100000 iterations on the spike-and-slab density with different dimensions ( $d = 2, 5, 10, 15$ ). Specifically, we plot the proportion of the Gibbs samples taken from the first component (the spike) for  $\mathbf{x}$  of different dimensions.


The mixing of the sampler for lower dimensions ( $d = 2$  and  $d = 5$ ) is reasonable, but we can see that for higher dimensions the sampler is extremely sticky and does not mix well at all. In the case of  $d = 10$ , the sampler only makes a couple movies between the spike component and the slab component during the run. In the case of  $d = 15$ , the sampler does not move into the second component at all during the entire 100000 sample run. We see this behavior more explicitly in Figure 3 where we plot the number of switches between components from the Gibbs sampling run for each dimension. Even higher dimensions ( $d = 20, d =$

25) were similar to the  $d = 15$  case.

**Figure 2:** Empirical Mixing Proportion after each iteration is calculated as the proportion of samples up to that iteration that are from the first component (the spike). Black lines indicate the empirical mixing proportion for the Gibbs sampler, gray lines indicate the true mixing proportion of 0.5



**Figure 3:** Number of component switches is the number of times during each Gibbs sampler run (out of 100000 iterations) that the algorithm moved from one mixture component to the other.





In high dimensions, we can estimate how often the Gibbs sampler will switch from one domain to the other. When the sampler is in one of the components, the variable  $x$  will mostly have a norm of  $\|\mathbf{x}\|_2^2 \approx d\sigma^2$  where  $\sigma = \sigma_0 = .05$  for component zero and  $\sigma = \sigma_1 = 3.0$  for component one. For a variable currently classified into the zero component, the probability of a switch is

$$P(I = 1 \mid \|\mathbf{x}\|_2^2 \approx d\sigma_0^2) \approx \left( \frac{\sigma_0\sqrt{e}}{\sigma_1} \right)^d$$


where the approximation holds if  $\sigma_0 \ll \sigma_1$ . Details of this calculation are given in Appendix C. This result suggests that the expected number of iterations until a switch is about  $(\sigma_1/\sigma_0\sqrt{e})^d = 36.4^d$ . This approximate switching time will be more accurate for large  $d$ .

### 3.2 Level-set Hit-and-run Sampling for Spike and Slab

Our LSHR1 sampling methodology was implemented on the same spike-and-slab distribution. Our sampling algorithm does not depend on a data augmentation scheme that moves between the two components. Rather, we start from the posterior mode and move outwards in convex level sets, while running a hit-and-run algorithm within each level set. As outlined in Section 2.4, each convex level set is defined by a threshold on the density that is slowly decreased in order to assure that each level set has a warm start.

Figure 4 shows the number of thresholds needed to fully explore the spike-and-slab density. Even for high dimensions ( $d = 10$  and  $d = 20$ ), we see that the spike-and-slab density is explored with a reasonable number of thresholds, and the linear trend suggests that our algorithm would scale well to even higher dimensions.


**Figure 4:** Number of thresholds needed for our level-set hit-and-run sampler on the spike-and-slab density as function of dimension.



The most striking comparison of the superior performance of our LSHR1 sampling strategy comes from examining the sampled values in any particular dimension relative to the true quantiles of the spike-and-slab mixture density. In Figure 5, we plot the true quantiles against the quantiles of the first dimension of sampled values from both the Gibbs sampler and the warm sampler.

We see that for low dimensions ( $d = 2$ ), samples from both the Gibbs sampler and warm sampler match the correct spike-and-slab distribution. However, for higher dimensions ( $d = 20$ ), the Gibbs sampler provides a grossly inaccurate picture of the true distribution, due to the fact that the Markov chain never escapes the spike component of the distribution. In contrast, the samples from our level-set hit-and-run sampler still provide a good match to the true distribution in high dimensions. We have checked even higher dimensions ( $d = 30$ ) and the level-set hit-and-run sampler still provides a good match to the true distribution.

**Figure 5:** Comparing level-set hit-and-run sampler and Gibbs sampler in terms of quantile-quantile plot of estimated vs. true spike-and-slab distribution. Top row is the low dimension ( $d = 2$ ) case, bottom row is the high dimension ( $d = 20$ ) case.



## 4 Empirical Study 2: Multivariate Normal

In our second empirical study, we compare our LSHR2 algorithm to the Gibbs sampler for obtaining samples from the multivariate normal density,

$$\boldsymbol{\theta} \sim \text{Normal}(\mathbf{0}, \boldsymbol{\Sigma}) \quad (2)$$

This is a superficial setting in the sense that it is easy to obtain samples  $\boldsymbol{\theta}$  directly from a multivariate normal density. However, this simple example is still illustrative of differences in performance between our LSHR2 level-set sampler and a component-wise Gibbs sampler.

We considered several choices of dimension  $d$  ( $d = 2, 10, 20, 30$ ) and also several specifications of the data variance  $\boldsymbol{\Sigma}$ . Specifically, we assume all diagonal

elements of  $\Sigma$  are one, but we consider several different values for off-diagonal elements  $\rho$  ( $\rho = 0, 0.5, 0.9, 0.99$ ).

We use our LSHR2 algorithm to obtain samples  $\theta$  by re-formulating (2) as the exponentially-tilted quasi-concave posterior density that results from a log-concave multivariate normal likelihood of  $y = \mathbf{0}|\theta, \Sigma$  and a bounded uniform prior density on  $\theta$  (i.e.  $\theta_k \sim \text{Unif}(-6, 6)$  for  $k = 1, \dots, d$ ). In this simple situation where our quasi-concave prior density is uniform, we only have to use one level set in our LSHR2 algorithm.

We also use a component-wise Gibbs sampler to obtain samples of  $\theta$  by iteratively sampling from the conditional normal densities of  $\theta_k|\theta_{(-k)}$ . Both the Gibbs sampler and LSHR2 sampler were run for one million iterations. In the case of the Gibbs sampler, the first 50000 iterations were discarded as burn-in.

In Figure 6, we examine the performance of our LSHR2 algorithm versus the Gibbs sampler when  $d = 2$  and both dimensions are independent (i.e.  $\rho = 0$ ). Specifically, we give a histogram of the samples for the first dimension  $\theta_1$  of  $\theta$  and the empirical autocorrelation function of the  $\theta_1$  samples for both Gibbs and LSHR1 algorithms.

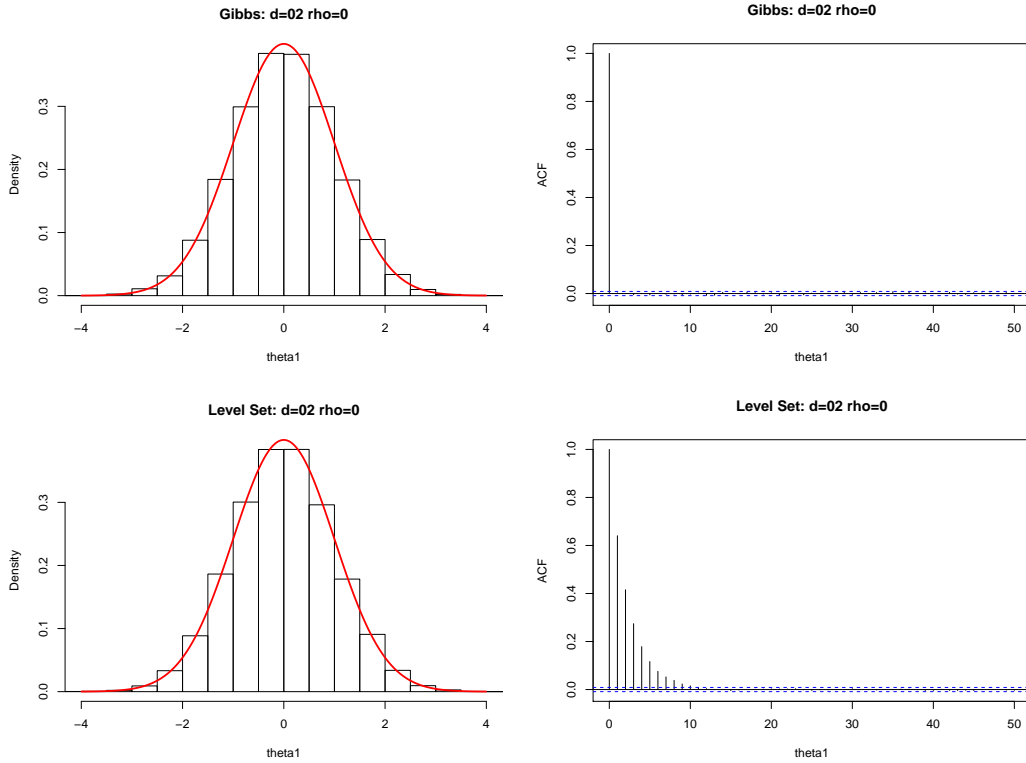
We see in Figure 6 that the samples of  $\theta_1$  from both algorithms are a close match to the true density. We also see that the Gibbs sampler has less autocorrelation than our level-set procedure in this setting where the true correlation between the dimensions of  $\theta$  is zero. Compare this observation with Figure 7, where we examine the performance of our LSHR2 algorithm versus the Gibbs sampler when  $d = 2$  but the correlation between the dimensions is high, i.e.  $\rho = 0.99$ .

When the dimensions of  $\theta$  are highly correlated, we see in Figure 6 that the performance of the Gibbs sampler is severely degraded, as evidenced by very

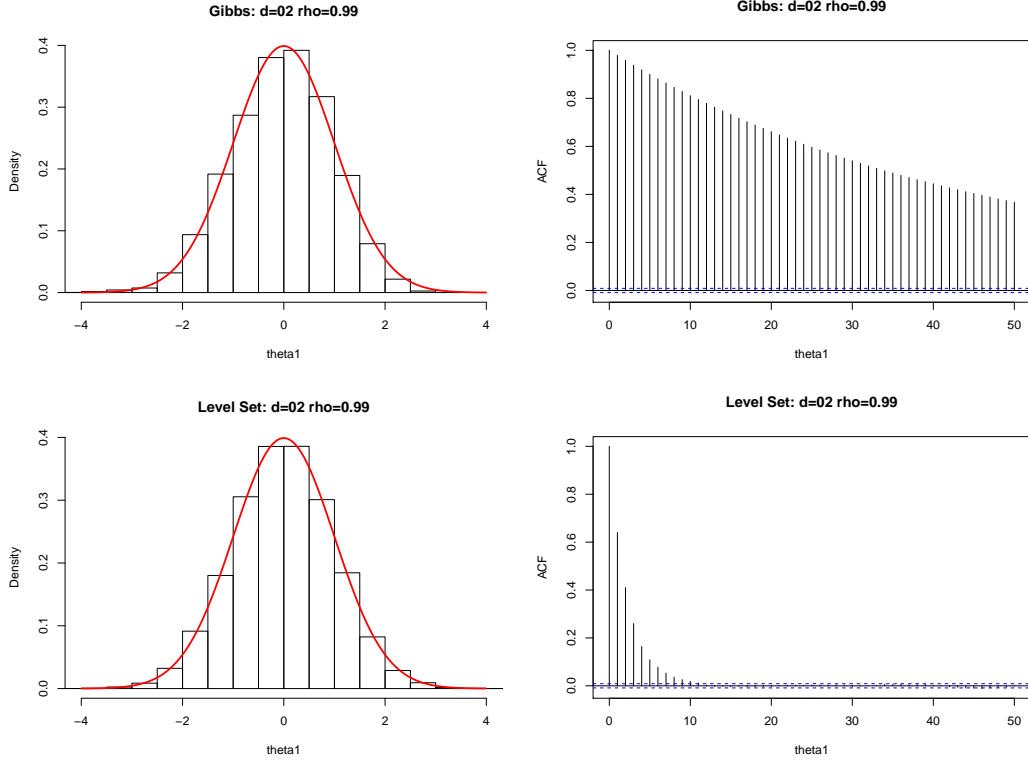
high autocorrelation of sampled values. In contrast, the autocorrelation of our LSHR2 algorithm is essentially unchanged between Figures 6 and 7, which indicates our methodology is robust to high dependence among the parameters  $\theta$ .

These same results were seen to an even greater extent when we compared the performance of the Gibbs sampler and LSHR2 algorithms in higher dimensions ( $d = 10, 20, 30$ ). Even in this artificial setting of a multivariate normal density with high correlation, we see that our level set methodology offers advantages over a component-wise Gibbs sampling strategy. In the next section, we examine a more complicated model where our LSHR2 algorithm also out-performs Gibbs sampling.

**Figure 6:** Comparison of samples of  $\theta$  from multivariate normal density with  $d = 2$  and  $\rho = 0$ . Top row is samples from Gibbs sampler. Bottom row is samples from LSHR2 algorithm. Left column is the histograms of the sampled values of  $\theta_1$  (red line indicates true density). Right column is the empirical autocorrelation functions of the sampled values of  $\theta_1$ .



**Figure 7:** Comparison of samples of  $\theta$  from multivariate normal density with  $d = 2$  and  $\rho = 0.99$ . Top row is samples from Gibbs sampler. Bottom row is samples from LSHR2 algorithm. Left column is the histograms of the sampled values of  $\theta_1$  (red line indicates true density). Right column is the empirical autocorrelation functions of the sampled values of  $\theta_1$ .



## 5 Empirical Study 3: Cauchy-normal model

In our third empirical comparison, we compare our LSHR2 algorithm to the Gibbs sampler on a Bayesian model consisting of a multivariate normal likelihood and a Cauchy prior density,


$$\begin{aligned} \mathbf{y}|\boldsymbol{\theta} &\sim \text{Normal}(\boldsymbol{\theta}, \sigma^2 \mathbf{I}) \\ \boldsymbol{\theta} &\sim \text{Cauchy}(\mathbf{0}, \mathbf{I}) \end{aligned} \tag{3}$$

where  $\mathbf{y}$  and  $\boldsymbol{\theta}$  are  $d$  dimensions and  $\sigma^2$  is fixed and known. As mentioned in Section 2.3, this combination of a log-concave density  $g(\mathbf{y}|\boldsymbol{\theta})$  and a quasi-

concave density for  $f(\boldsymbol{\theta})$  gives an *exponentially-tilted* quasi-concave posterior density  $h(\boldsymbol{\theta}|\mathbf{y})$ .

For some combinations of  $\sigma^2$  and observed  $\mathbf{y}$  values, the posterior density  $h(\boldsymbol{\theta}|\mathbf{y})$  is multi-modal. Figure 8 gives examples of this multi-modality in one and two dimensions. The one-dimensional  $h(\boldsymbol{\theta}|\mathbf{y})$  in Figure 8a has  $y = 10$  and  $\sigma^2 = 10.84$ , whereas the two-dimensional  $h(\boldsymbol{\theta}|\mathbf{y})$  in Figure 8b has  $y = (10, 10)$  and  $\sigma^2 = 12.57$ . In higher dimensions  $d$ , we used  $y = (10, \dots, 10)$  and  $\sigma^2 = d \cdot 10^2 / ((d + 1) \log(1 + d \cdot 10^2))$ , which is a formula derived from equating the density value at the prior mode to the density at the likelihood mode.

**Figure 8:** Posterior density  $h(\boldsymbol{\theta}|\mathbf{y} = 10)$  for Cauchy-normal model in (a) one and (b) two dimensions.



In higher dimensions ( $d \geq 3$ ), it is difficult to evaluate (or sample from) the true posterior density  $h(\boldsymbol{\theta}|\mathbf{y})$  which we need in order to have a gold standard for comparison between the Gibbs sampler and LSHR2 algorithm. Fortunately, for this simple model, there is a rotation procedure that we detail in our supplementary materials that allows us to accurately estimate the true posterior density  $h(\boldsymbol{\theta}|\mathbf{y})$ .

## 5.1 Gibbs Sampling for Cauchy Normal Model

The Cauchy-normal model (3) can be estimated via Gibbs sampling by augmenting the parameter space with an extra scale parameter  $\tau^2$ ,

$$\begin{aligned} \mathbf{y}|\boldsymbol{\theta} &\sim \text{Normal}(\boldsymbol{\theta}, \sigma^2 \mathbf{I}) \\ \boldsymbol{\theta}|\tau^2 &\sim \text{Normal}(\mathbf{0}, \tau^2 \mathbf{I}) \\ (\tau^2)^{-1} &\sim \text{Gamma}\left(\frac{1}{2}, \frac{1}{2}\right) \end{aligned} \tag{4}$$

Marginalizing over  $\tau^2$  gives us the same  $\text{Cauchy}(\mathbf{0}, \mathbf{I})$  prior for  $\boldsymbol{\theta}$ . The Gibbs sampler iterates between sampling from the conditional distribution of  $\boldsymbol{\theta}|\tau^2, \mathbf{y}$ :

$$\boldsymbol{\theta}|\tau^2, \mathbf{y} \sim \text{Normal}\left(\frac{\frac{1}{\sigma^2}\mathbf{y}}{\frac{1}{\sigma^2} + \frac{1}{\tau^2}}, \frac{1}{\frac{1}{\sigma^2} + \frac{1}{\tau^2}}\mathbf{I}\right) \tag{5}$$

and the conditional distribution of  $\tau^2|\boldsymbol{\theta}, \mathbf{y}$ :

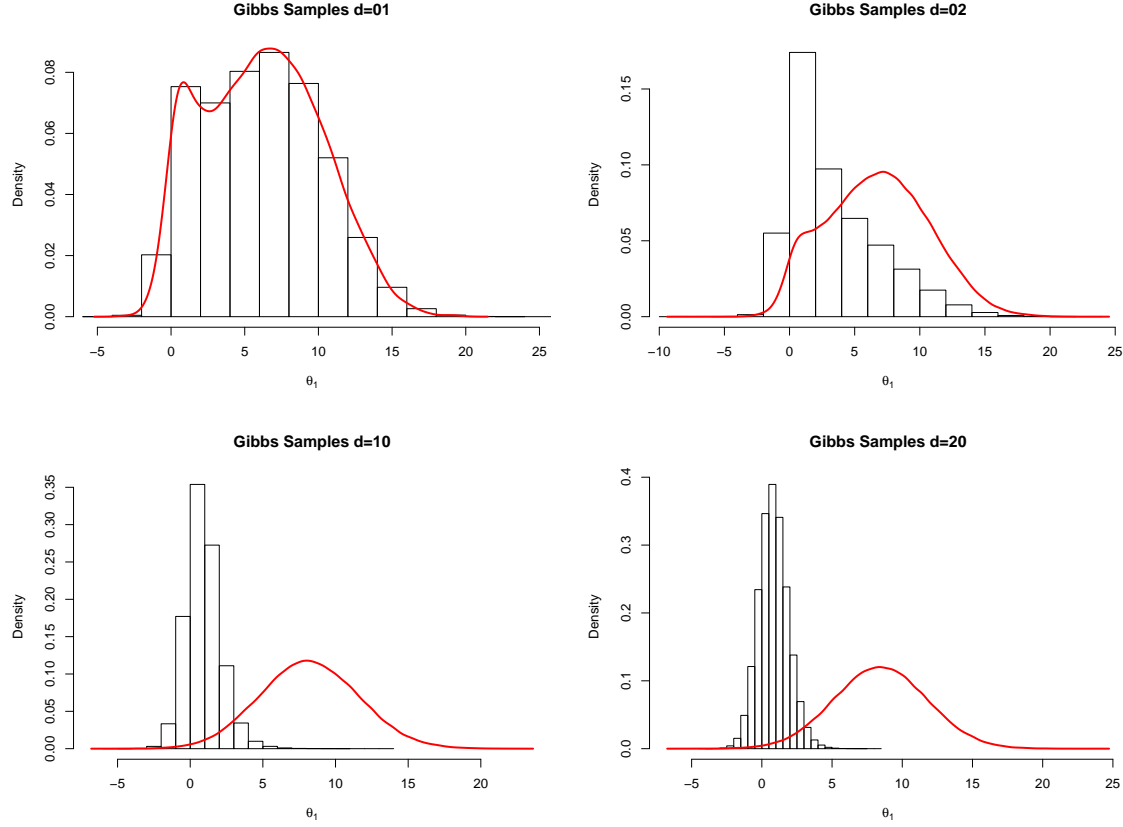
$$(\tau^2)^{-1}|\boldsymbol{\theta}, \mathbf{y} \sim \text{Gamma}\left(\frac{d+1}{2}, \frac{d + \boldsymbol{\theta}^T \boldsymbol{\theta}}{2}\right) \tag{6}$$

For several different dimensions  $d$ , we ran this Gibbs sampler for one million iterations, with the first 500000 iterations discarded at burn-in. In Figure 9, we compare the posterior samples for the first dimension  $\theta_1$  from our Gibbs sampler to the true posterior distribution for  $d = 1, 2, 10$  and 20.

We see that the Gibbs sampler performs poorly at exploring the full posterior space in dimensions greater than one. Even in two dimensions, the Gibbs sampler struggles to fully explore both modes of the posterior distribution and this problem is exacerbated in dimensions higher than two.



**Figure 9:** Posterior samples of the first dimension  $\theta_1$  of  $\theta$  from Gibbs sampler for Cauchy-normal model in different dimensions  $d$ . Red curve in each plot represents the true posterior density.




## 5.2 LSHR2 Level Set Sampling for Cauchy-Normal Model

Our LSHR2 level set sampling methodology was implemented on the same Cauchy-normal model. We start from the prior mode and move outwards in convex level sets, while running an exponentially-tilted hit-and-run algorithm within each level set in order to get samples from the posterior distribution. As outlined in Section 2.4, each convex level set is defined by a threshold on the density that is slowly decreased in order to assure that each level set has a warm start.

Figure 10 shows the number of thresholds needed to fully explore the posterior

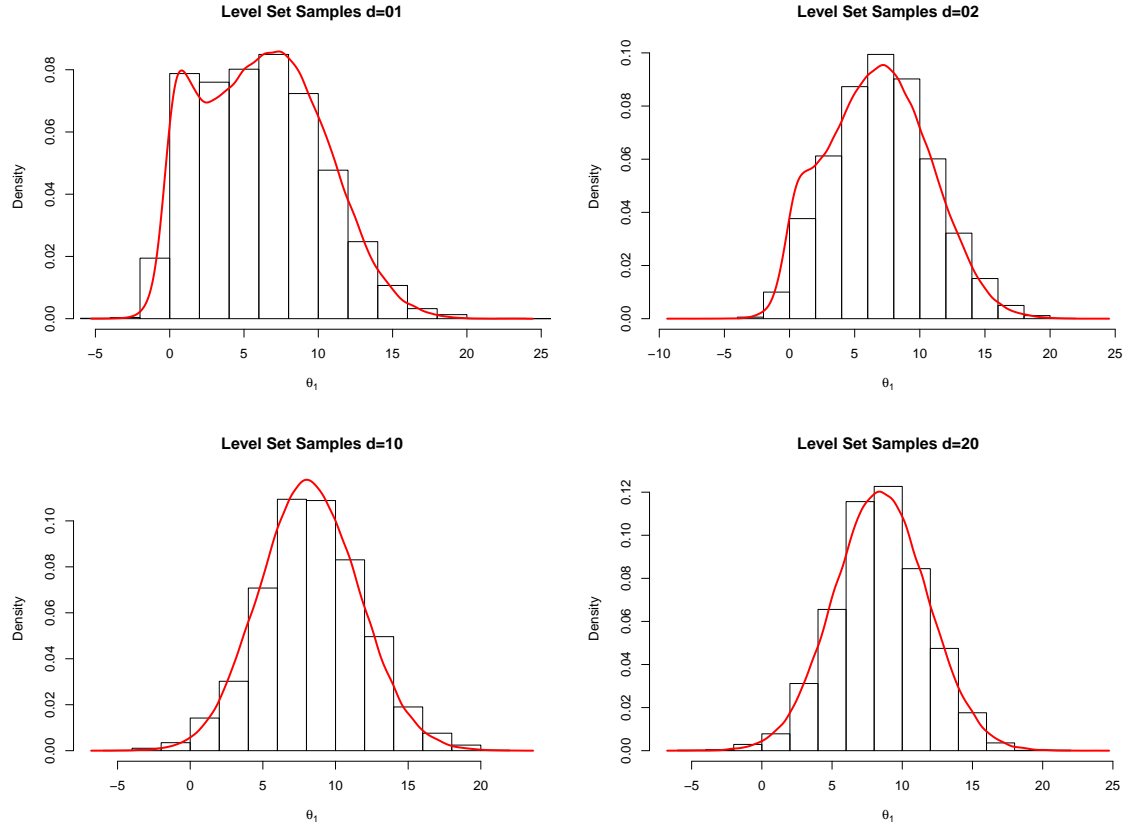
distribution of the Cauchy-normal model for different number of dimensions. Even for high dimensions ( $d = 10$  and  $d = 20$ ), we see that the posterior density is explored with a reasonable number of thresholds, and the approximately linear trend suggests that our algorithm would scale well to even higher dimensions.

**Figure 10:** Number of thresholds needed for our level-set hit-and-run sampler on the posterior density of the Cauchy-normal model as function of dimension.



In Figure 11, we compare the posterior samples for the first dimension  $\theta_1$  from our LSHR2 sampler to the true posterior distribution for  $d = 1, 2, 10$  and  $20$ . Our LSHR2 level set sampler (with exponential tilting) samples closely match the true posterior density in both low and high-dimensional cases. Comparing Figures 9 and Figures 11 clearly suggests that our level set hit-and-run methodology gives more accurate samples from the Cauchy-normal model in dimensions higher than one.

**Figure 11:** Posterior samples of the first dimension  $\theta_1$  of  $\theta$  from LSHR2 sampler for Cauchy-normal model in different dimensions  $d$ . Red curve in each plot represents the true posterior density.



## 6 Discussion

In this paper, we have developed a general sampling strategy based on dividing a density into a series of level sets, and running the hit-and-run algorithm within each level set. Our basic level-set hit-and-run sampler (LSHR1) can be applied to the broad class of pseudo-concave densities, which includes most distributions used in applied statistical modeling. We illustrate our LSHR1 sampler on spike-and-slab density (Section 3), where our procedure performs much better in high dimensions than the usual Gibbs sampler.

We also extend our sampling methodology to an exponentially-tilted level-set hit-and-run sampler (LSHR2). Our LSHR2 algorithm can be applied to log-concave densities, which we illustrate with a simple multivariate normal example in Section 4. In this setting, our LSHR2 algorithm is more effective than Gibbs sampling when there is high dependence between parameters.

Our LSHR2 can also be applied to exponentially-tilted quasi-concave densities, which arise frequently in Bayesian models as the posterior distribution formed from the combination of a log-concave likelihood and quasi-concave prior density. We illustrate our LSHR2 sampler on the posterior density from a Cauchy normal model (Section 5), which can exhibit multi-modality under certain parameter settings. Our procedure is able to sample effectively from this multi-modal posterior density even in high dimensions, where the Gibbs sampler performs poorly.

## A The LSHR1 Algorithm

In this appendix, we give pseudo-code for our LSHR1 level-set hit-and-run sampler described in Section 2.2. We must pre-specify a minimum value  $K$  of the probability density  $f(\cdot)$  as our stopping threshold. The first threshold  $t_1$  must be pre-specified, and we arbitrarily use  $t_1 = 0.95 * f(\mathbf{x}_{\max})$ . Finally, the number of hit-and-run iterations per level set,  $m$ , must be chosen. In our spike-and-slab example, we set the number of hit-and-run iterations per level set to  $m = 1000$  based on the evaluation presented in our supplementary materials.

---

### Algorithm 1 LSHR1: Level-set Hit-and-run Sampler

---

```

1: initialize  $\mathbf{x} = \mathbf{x}_{\max}$  and  $\Sigma_1 = \mathbf{I}$ 
2: initialize  $t_1$  and set  $k = 1$ 
3: define current level set  $C_1$  as  $\{\mathbf{x} : f(\mathbf{x}) > t_1\}$ 
4: for  $m$  iterations do
5:   sample random direction  $\mathbf{d}$ 
6:   calculate boundaries  $\mathbf{a}$  and  $\mathbf{b}$  of  $C_1$  along direction  $\mathbf{d}$ 
7:   sample new  $\mathbf{x}$  uniformly between  $\mathbf{a}$  and  $\mathbf{b}$ 
8:   store all samples  $\mathbf{x}$  in  $\{\mathbf{x}\}_1$ 
9: while  $t_k > K$  do
10:  propose new  $t_{\text{prop}} < t_k$ 
11:  define proposed level set  $C_{\text{prop}}$  as  $\{\mathbf{x} : f(\mathbf{x}) > t_{\text{prop}}\}$ 
12:  for  $m$  iterations do
13:    sample random direction  $\mathbf{d}^*$  and set  $\mathbf{d} = \Sigma_k^{1/2} \mathbf{d}^*$ 
14:    calculate boundaries  $\mathbf{a}$  and  $\mathbf{b}$  of  $C_1$  along direction  $\mathbf{d}$ 
15:    sample new  $\mathbf{x}$  uniformly between  $\mathbf{a}$  and  $\mathbf{b}$ 
16:    if  $f(\mathbf{x}) > t_k$  then  $\hat{R}_{k:\text{prop}} = \hat{R}_{k:\text{prop}} + 1$ 
17:   $\hat{R}_{k:\text{prop}} = \hat{R}_{k:\text{prop}} / m$ 
18:  if  $0.55 \leq \hat{R}_{1:\text{prop}} \leq 0.8$  then
19:     $k = k + 1$ 
20:     $t_k = t_{\text{prop}}$ ;  $C_k = C_{\text{prop}}$ ;  $\hat{R}_{k:k+1} = \hat{R}_{k:\text{prop}}$ ;  $\Sigma_{k+1} = \text{Cov}(\mathbf{x})$ 
21:    store all samples  $\mathbf{x}$  in  $\{\mathbf{x}\}_k$ 
22: define  $n$  = number of level sets
23: for  $i = 1, \dots, n$  do
24:   $q_i = (t_{i-1} - t_i) \times \prod_{k=i}^n \hat{R}_{k:k+1}$ 
25:  where  $t_0$  is the maximum of  $f(\cdot)$  and  $\hat{R}_{n:n+1} = 1$ 
26: for  $l = 1, \dots, L$  do
27:  Sample level set  $s$  w.p.  $p_s = q_s / \sum q_s$ 
28:  Sample point  $\mathbf{x}_l$  randomly from level set collection  $\{\mathbf{x}\}_s$ 

```

---

When considering the proposal for level set  $k + 1$ , the initial  $t_{\text{prop}}$  is set to  $t_k - (t_{k-1} - t_k)$ , so that our proposal moves the same distance in  $t$  as the previous successful move. If this initial  $t_{\text{prop}}$  is rejected for not being warm enough, we choose a new proposal  $t_{\text{prop}}$  as the midpoint of the old proposal and  $t_k$ , and so on.

## B The LSHR2 Algorithm

In this appendix, we give pseudo-code for LSHR2, our exponentially-weighted level-set hit-and-run sampler procedure described in Section 2.4. We must pre-specify a minimum value  $K$  of the probability density  $f(\cdot)$  as our stopping threshold. The first threshold  $t_1$  must be pre-specified, and we arbitrarily use  $t_1 = 0.95 * f(\theta_{\max})$  where  $\theta_{\max}$  is the mode of  $f(\theta)$ . We set the number of hit-and-run iterations per level set to  $m = 1000$ .

---

### Algorithm 2 LSHR2: Exponentially-tilted Level-set Hit-and-run Sampler

---

```

1: initialize  $\theta = \theta_{\max}$  and  $\Sigma_1 = \mathbf{I}$ 
2: initialize  $p = \log g(\theta_{\max})$  and  $\theta^* = (\theta, p)$ 
3: initialize  $t_1$  and set  $k = 1$ 
4: define current level set  $D_1$  as  $\{\theta^* : f(\theta) \geq t_1, p < \log g(\mathbf{y}|\theta)\}$ 
5: for  $m$  iterations do
6:   sample random direction  $\mathbf{d}$  in  $(d + 1)$  dimensions
7:   calculate boundaries  $\mathbf{a}$  and  $\mathbf{b}$  of  $D_1$  along direction  $\mathbf{d}$ 
8:   sample new  $\theta^*$  proportional to  $\exp(p)$  from line segment between  $\mathbf{a}$  and  $\mathbf{b}$ 
9:   store all samples  $\theta^*$  in  $\{\theta^*\}_1$ 
10: while  $t_k > K$  do
11:   propose new  $t_{\text{prop}} < t_k$ 
12:   define proposed level set  $D_{\text{prop}}$  as  $\{\theta^* : f(\theta) \geq t_{\text{prop}}, p < \log g(\mathbf{y}|\theta)\}$ 
13:   for  $m$  iterations do
14:     sample random direction  $\mathbf{d}^*$  in  $(d + 1)$  dimensions and set  $\mathbf{d} = \Sigma_k^{1/2} \mathbf{d}^*$ 
15:     calculate boundaries  $\mathbf{a}$  and  $\mathbf{b}$  of  $D_{\text{prop}}$  along direction  $\mathbf{d}$ 
16:     sample new  $\theta^*$  proportional to  $\exp(p)$  from line segment between  $\mathbf{a}$  and  $\mathbf{b}$ 
17:     if new  $\theta^* \in D_k$  then  $\hat{R}_{k:\text{prop}} = \hat{R}_{k:\text{prop}} + 1$ 
18:    $\hat{R}_{k:\text{prop}} = \hat{R}_{k:\text{prop}}/m$ 
19:   if  $0.55 \leq \hat{R}_{1:\text{prop}} \leq 0.8$  then
20:      $k = k + 1$ 
21:      $t_k = t_{\text{prop}}; D_k = D_{\text{prop}}; \hat{R}_{k:k+1} = \hat{R}_{k:\text{prop}}; \Sigma_{k+1} = \text{Cov}(\mathbf{x})$ 
22:     store all samples  $\theta^*$  in  $\{\theta^*\}_k$ 
23: define  $n$  = number of level sets
24: for  $i = 1, \dots, n$  do
25:    $q_i = (t_{i-1} - t_i) \times \prod_{k=i}^n \hat{R}_{k:k+1}$ 
26:   where  $t_0$  is the maximum of  $f(\cdot)$  and  $\hat{R}_{n:n+1} = 1$ 
27: for  $l = 1, \dots, L$  do
28:   Sample level set  $s$  w.p.  $p_s = q_s / \sum q_s$ 
29:   Sample point  $\theta_l^*$  randomly from level set collection  $\{\theta^*\}_s$ 

```

---

When considering the proposal for level set  $k + 1$ , the initial  $t_{\text{prop}}$  is set to  $t_k - (t_{k-1} - t_k)$ , so that our proposal moves the same distance in  $t$  as the previous successful move. If this initial  $t_{\text{prop}}$  is rejected for not being warm enough, we choose a new proposal  $t_{\text{prop}}$  as the midpoint of the old proposal and  $t_k$ , and so on.

## C Gibbs Switching Calculation

For a variable in the zero component, with norm  $\|\mathbf{x}\|_2^2 \approx d\sigma_0^2$ , the probability of a switch is

$$\begin{aligned}
 P(I = 1 \mid \|\mathbf{x}\|_2^2 \approx d\sigma_0^2) &= \frac{\sigma_1^{-d} \exp(-\|\mathbf{x}\|_2^2/2\sigma_1^2)}{\sigma_1^{-d} \exp(-\|\mathbf{x}\|_2^2/2\sigma_1^2) + \sigma_0^{-d} \exp(-\|\mathbf{x}\|_2^2/2\sigma_0^2)} \\
 &= \frac{\sigma_1^{-d} \exp(-d\sigma_0^2/2\sigma_1^2)}{\sigma_1^{-d} \exp(-d\sigma_0^2/2\sigma_1^2) + \sigma_0^{-d} \exp(-d/2)} \\
 &= \frac{\left(\frac{\sigma_0}{\sigma_1}\right)^d \exp\left(\frac{d}{2}\left(1 - \frac{\sigma_0^2}{\sigma_1^2}\right)\right)}{\left(\frac{\sigma_0}{\sigma_1}\right)^d \exp\left(\frac{d}{2}\left(1 - \frac{\sigma_0^2}{\sigma_1^2}\right)\right) + 1} \\
 &\approx \frac{\left(\frac{\sigma_0}{\sigma_1}\right)^d \exp\left(\frac{d}{2}\right)}{\left(\frac{\sigma_0}{\sigma_1}\right)^d \exp\left(\frac{d}{2}\right) + 1}
 \end{aligned}$$

since  $\sigma_0 \ll \sigma_1$ . Now, as  $d$  increases,  $(\sigma_0/\sigma_1)^d$  goes to zero, and so

$$\frac{\left(\frac{\sigma_0}{\sigma_1}\right)^d \exp\left(\frac{d}{2}\right)}{\left(\frac{\sigma_0}{\sigma_1}\right)^d \exp\left(\frac{d}{2}\right) + 1} \approx \left(\frac{\sigma_0}{\sigma_1}\right)^d \exp\left(\frac{d}{2}\right) = \left(\frac{\sigma_0 \sqrt{e}}{\sigma_1}\right)^d$$

## References

- Bagnoli, M. and Bergstrom, T. (2005). Log-concave probability and its applications. *Economic Theory* **26**, 445–469.
- Chen, M.-H. and Schmeiser, B. W. (1996). General hit-and-run monte carlo sampling for evaluating multidimensional integrals. *Operations Research Letters* **19**, 161–169.
- Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transaction on Pattern Analysis and Machine Intelligence* **6**, 721–741.
- George, E. I. and McCulloch, R. E. (1997). Approaches for bayesian variable selection. *Statistica Sinica* **7**, 339–373.
- Hastings, W. (1970). Monte carlo sampling methods using markov chains and their applications. *Biometrika* **57**, 97–109.
- Lovasz, L. (1999). Hit-and-run mixes fast. *Mathematical Programming* **86**, 443–461.
- Vempala, S. (2005). Geometric random walks: A survey. *Combinatorial and Computational Geometry* **52**, 573–612.